



FIFE Workshop 2014

What's new with CVMFS



TOPIC FOR WHAT'S NEW IN FIFE

DAVE DYKSTRA

Major upgrades

2

- Since last year's workshop, cvmfs 2.1 clients have been installed at a large number of OSG sites
- About a dozen VOs are active in the OSG OASIS CVMFS distribution
- As the number of files grew, the CVMFS catalog became enormous and unwieldy
 - There are now separate catalogs for each VO in the oasis repository, and subcatalogs within each VO, generally one for each package release
 - Performance is greatly improved

Major upgrades

3

- The CERN/OSG Stratum 1s were all upgraded to CVMFS 2.1
 - Much easier to manage
 - Able to support 2.1 repositories which include many important features
 - ✦ Multiple repositories on the same repository server
 - ✦ Chunking of large files into smaller pieces
 - ✦ Transaction-based updates
 - ✦ Avoids needing an entire unpacked copy of the source files; uses a cvmfs client to view all the existing files in the compressed repo
- Multiple new client versions have come out over the last year with new features and bug fixes

Problem with slow updates understood

4

- CVMFS updates should be available within a couple of hours
 - Rarely it can be a little longer because stratum 1 busy updating other repositories
- There are however occasions where updates appear to be many hours behind
 - In those cases the cvmfs client is taking one minute to flush buffers and let's the application continue
 - Workaround: run `/cvmfs/grid.cern.ch/util/cvmfs-uptodate` in pilot or beginning of job
 - Permanent fixes being considered for future client release

OSG repositories hosted at Fermilab


5

- OSG now has the ability to distribute to the grid CVMFS repositories that are hosted at sites
 - Can mount on worker nodes with no change to them if they were set up as instructed
 - /cvmfs/fermilab.opensciencegrid.org is being set up as the first one, hosted here, a 2.1 repository
 - ✦ Evolution of /cvmfs/oasis.opensciencegrid.org/fermilab
 - ✦ Contains common UPS/UPD products and larsoft
 - Expect /cvmfs/nova.opensciencegrid.org to be next
 - Should become almost as easy to set up as fnal.gov repository
 - Access will be through ssh & Kerberos, not gsissh

Using CVMFS for auxiliary data files

6

- CVMFS is designed for distributing code, but many experiments have a need to distribute some data to a lot of nodes as well
- Prime example: neutrino “flux” files for the GENIE application
 - Measured a case of 128 jobs each reading 1.8GB of 140MB flux files, out of a 15GB dataset
 - Jobs ran on 87 different nodes
 - 4 2Gbit/s squids were used, but average transfer rate mostly limited to disk speeds ($\sim 4 \times 30\text{MB/s}$, $\sim 1.2\text{Gbit/s}$)



Using CVMFS for auxiliary data files

8

- Several different possible solutions have been proposed
- One we're testing is using the CVMFS 2.1.17 client “Alien Cache” feature
 - Separate repositories for this type of data would be created, and configured in clients to store on local Storage Element (e.g. dCache, Lustre, HDFS, EOS)
 - Cache would be shared between the worker nodes
 - Files would get automatically transferred from central repo on demand
 - Not a solution for all types of data files

Coming up

9

- Sharing of repositories between OSG & EGI coming soon
 - Will require a config rpm upgrade on worker nodes
- Next cvmfs release will support distributing CVMFS configs in a CVMFS repository, so additional trusted distributions can be added centrally
 - Includes new method to support automated ordering of stratum 1 servers by distance

FIFE CVMFS Documentation

10

- [https://cdcvs.fnal.gov/redmine/projects/fife/wiki/Introduction_to_FIFE_and_Component_Services#OASISCV MFS](https://cdcvs.fnal.gov/redmine/projects/fife/wiki/Introduction_to_FIFE_and_Component_Services#OASISCV_MFS)